

What is claimed is:

- 1        1. A method of determining a status of a peer protocol initiated on a plurality  
2        of members of a group in a clustered computer system, the method comprising:
  - 3                (a) locally tracking protocol progress information within each member  
4                of the group; and
  - 5                (b) responding to a query directed to a selected member of the group  
6                by providing the protocol progress information locally tracked by the selected  
7                member.
- 1        2. The method of claim 1, wherein locally tracking protocol progress  
2        information includes tracking, within a first member of the group, acknowledgment  
3        (ACK) messages directed to the first member by each other member of the group.
- 1        3. The method of claim 1, wherein locally tracking protocol progress  
2        information includes:
  - 3                (a) tracking, within a first member of the group, a current  
4                acknowledgment (ACK) round for the first member, the current ACK round  
5                associated with a current peer protocol being processed by the first member;  
6                and
  - 7                (b) tracking, within the first member, a last ACK round received  
8                parameter associated with each other member of the group, the last ACK  
9                round received parameter for each other member identifying a peer protocol  
10              associated with a last received ACK message from such other member.
- 1        4. The method of claim 3, wherein locally tracking protocol progress  
2        information further includes updating the current ACK round for the first member in  
3        response to receipt of ACK messages for the current peer protocol from all other  
4        members of the group.
- 1        5. The method of claim 1, wherein locally tracking protocol progress  
2        information includes updating the protocol progress information for a first member of

3 the group in response to receipt of an acknowledgment (ACK) message directed to the  
4 first member.

1 6. The method of claim 1, further comprising:

2 (a) waiting on a resource required by a protocol being processed on  
3 the selected member; and  
4 (b) monitoring for receipt of the query by the selected member while  
5 waiting on the resource.

1 7. The method of claim 6, wherein the protocol is a peer protocol, and  
2 wherein waiting on the resource includes waiting for receipt of an acknowledgment  
3 (ACK) message directed to the selected member.

1 8. The method of claim 6, wherein the protocol is a local protocol, and  
2 wherein waiting on the resource includes waiting on a local resource requested by the  
3 selected member.

1 9. The method of claim 8, wherein the local resource is selected from the  
2 group consisting of a lock and a creation of a new job.

1 10. The method of claim 6, wherein waiting on the resource includes waiting  
2 for receipt of a message by a local message queue for the selected member, and  
3 wherein monitoring for receipt of the query includes monitoring the local message  
4 queue for receipt of a query message.

1 11. The method of claim 1, wherein locally tracking protocol progress  
2 information within each member of the group includes locally tracking within the  
3 selected member protocol progress information associated with at least one other  
4 member in the group.

1           12. The method of claim 1, wherein locally tracking protocol progress  
2 information within each member of the group includes locally tracking within the  
3 selected member protocol progress information associated with all other members in  
4 the group.

1           13. The method of claim 1, wherein locally tracking protocol progress  
2 information within each member of the group includes locally tracking within each  
3 member protocol progress information associated with each other member in the  
4 group.

IBM ROC9-2000-0125-US1  
WH&E IBM/151

1       14. An apparatus, comprising:  
2           (a) a memory; and  
3           (b) a program resident in the memory, the program configured to  
4       determine a status of a peer protocol initiated on a plurality of members of a  
5       group in a clustered computer system by locally tracking protocol progress  
6       information within at least one member of the group, and providing the  
7       protocol progress information locally tracked by a member of the group in  
8       response to a query directed to such member.

1       15. The apparatus of claim 14, wherein the program is configured to locally  
2       track protocol progress information by tracking, within a first member of the group,  
3       acknowledgment (ACK) messages directed to the first member by each other member  
4       of the group.

1       16. The apparatus of claim 14, wherein the program is configured to locally  
2       track protocol progress information by tracking, within a first member of the group, a  
3       current acknowledgment (ACK) round for the first member, and tracking, within the  
4       first member, a last ACK round received parameter associated with each other  
5       member of the group, wherein the current ACK round is associated with a current  
6       peer protocol being processed by the first member, and wherein the last ACK round  
7       received parameter for each other member identifies a peer protocol associated with a  
8       last received ACK message from such other member.

1       17. The apparatus of claim 14, wherein the program is further configured to  
2       wait on a resource required by a protocol being processed on the selected member,  
3       and monitor for receipt of the query by the selected member while waiting on the  
4       resource.

1       18. The apparatus of claim 17, wherein the protocol is a peer protocol, and  
2       wherein the program is configured to wait on the resource by waiting for receipt of an  
3       acknowledgment (ACK) message directed to the selected member.

1           19. The apparatus of claim 17, wherein the protocol is a local protocol, and  
2        wherein the program is configured to wait on the resource by waiting on a local  
3        resource requested by the selected member.

1           20. The apparatus of claim 17, wherein the program is configured to locally  
2       track protocol progress information by locally tracking within a first member protocol  
3       progress information associated with at least one other member in the group.

1           21. The apparatus of claim 17, wherein the program is configured to locally  
2 track protocol progress information by locally tracking within each member protocol  
3 progress information associated with each other member in the group.

1           22. A clustered computer system, comprising:

2                   (a) a plurality of nodes coupled to one another over a network;

3                   (b) a plurality of member jobs defining a group and configured to be

4                   executed by at least one of the plurality of nodes; and

5                   (c) a program configured to be executed by at least one of the plurality

6                   of nodes to determine a status of a peer protocol initiated on the plurality of

7                   members by locally tracking protocol progress information within at least one

8                   member of the group, and providing the protocol progress information locally

9                   tracked by a member of the group in response to a query directed to such

10                  member.

1        23. A program product, comprising:

2            (a) a program configured to determine a status of a peer protocol  
3            initiated on a plurality of members of a group in a clustered computer system  
4            by locally tracking protocol progress information within at least one member  
5            of the group, and providing the protocol progress information locally tracked  
6            by a member of the group in response to a query directed to such member; and  
7

          (b) a signal bearing medium bearing the program.

1        24. The program product of claim 23, wherein the signal bearing medium  
2        includes at least one of a recordable medium and a transmission medium.

2025 RELEASE UNDER E.O. 14176

1        25. An apparatus, comprising:

2            (a) a memory; and

3            (b) a program, resident in the memory, the program configured to

4 monitor for receipt of a query message by a member of a group in a clustered

5 computer system while a current protocol for the member is waiting on a

6 resource, the program further configured to output protocol status information

7 in response to receipt of the query message.

1        26. The apparatus of claim 25, wherein the resource is selected from the

2 group consisting of a local resource and an acknowledgment (ACK) message.

IBM 6949 6950 6951 6952 6953 6954 6955 6956 6957 6958 6959 6960 6961 6962 6963 6964 6965 6966 6967 6968 6969 6970 6971 6972 6973 6974 6975 6976 6977 6978 6979 6980 6981 6982 6983 6984 6985 6986 6987 6988 6989 6990 6991 6992 6993 6994 6995 6996 6997 6998 6999